NIF Design Optimization

W. Williams
C. Orth
R. Sacks
J. Lawson
K. Jancaitis
J. Trenholme
S. Haney
J. Auerbach
M. Henesian
P. Renard

Introduction

There are many potential designs for the proposed National Ignition Facility (NIF) that would be expected to meet the machine's functional requirements, but at widely varying costs. Given the size of this project, an effort has been made to assure that we are, indeed, designing the least expensive option possible.

During the previous three years of NIF design, we have considerably refined computer codes used for such optimization analysis. During the conceptual design phase of NIF, we used a simplified code, culminating in the configuration presented in the *Conceptual Design Report* (CDR)[1] published in May 1994. For the NIF Advanced Conceptual Design (ACD), we used a grouping of several codes that allowed design optimization coupled with more sophisticated propagation models. In this article, we review some of the assumptions and modeling procedures used in both sets of computer codes. We also summarize our results from the ACD optimization activity that resulted in the current (Title I) NIF design.

Computer Codes Used in NIF Optimization

Two computer codes were primarily used in NIF optimization: CHAINOP and PROPSUITE. The multiplatform code CHAINOP is the original optimization tool used during the CDR design process for NIF. CHAINOP's goal is to combine models of laser chain layout, cost, and performance in a way amenable to speedy analysis of several thousand potential system designs, identifying those least expensive configurations that fulfill functional requirements. (A laser "chain" refers to the sequence and spacing of optical components in one beam of a large laser system; there are 192 laser chains in NIF.) CHAINOP permits certain kinds of constraints, such as damage fluence limits, B-integral limits, and irradiance modulation limits, that guide the optimization into safe laser designs via simple, often heuristic models. With a ray-trace model for tracking beam size changes and single-value loss and saturating gain models for energetics, CHAINOP reveals little specific detail of spatial and temporal characteristics of pulse propagation. These apparent simplifications are somewhat deceptive, however. Many of the heuristic models contained in CHAINOP are the result of detailed studies of physical processes, and allow the code to predict reliably laser designs that are consistent with the more robust modeling tools that were subsequently developed.

As our confidence grew in the modeling efforts for individual processes like diffraction propagation, frequency conversion, amplifier pumping, and optics damage, we were able to assemble individual computer codes into a more sophisticated optimization package called PROPSUITE. The backbone of PROPSUITE is a beam propagation code called PROP92 that uses Fourier methods to follow temporal and spatial details of a laser pulse as it traverses an optical chain. Among other things, this code includes effects of paraxial, nonlinear beam diffraction, beam movement for vignetting purposes, energy clipping at pinholes, and clipping of the beam on a hard aperture. It also allows for inclusion of detailed phase representations of optics surfaces. For a detailed discussion of PROP92, see p. 207 of this *Quarterly*. During optimizations, the laser was modeled along the direction of beam propagation, and in one transverse dimension—in this case, along the horizontal axis. The spatial resolution considered was up to ~32,000 points over a 50-cm aperture, sufficient to resolve the fine spatial features of the optics surfaces. Simulations involving both transverse dimensions are too slow to put in an optimization loop; however, designs from the one-dimensionsal (1-D) optimizations have been further analyzed and confirmed in two-dimensional (2-D) simulations.

The PROPSUITE collection also includes frequency-conversion computer codes with either plane-wave models or a 1-D spatio-temporal model of conversion (THGFT02 and THGXTZ002). Specifics of amplifier flashlamp pumping are modeled by a 2-D optical ray-trace computer code. The thermo-mechanical effects that cause pump-induced distortions are handled with public-domain software TOPAZ and NIKE3D. PROPSUITE also used the SUPERCODE shell, exercised extensively for tokamak systems studies, to perform multivariable optimizations. Additional codes were developed to do system layout calculations, based on rules in CHAINOP, and to tie the codes together. For optimizations, calculations of several thousand possible designs were done on a cluster of 28 workstations using parallel processing with the public domain software PVM.

In the following sections we discuss some of the basic modeling concepts contained in both CHAINOP and PROPSUITE.

Chain Layout

Both CHAINOP and PROPSUITE were designed around the NIF multipass architecture of a relayed, mirrored cavity containing one or two amplifiers, a spatial filter, and a Pockels cell switch/polarizer. The cavity is followed by a booster amplifier, transport spatial filter, transport mirrors, frequency converter, final focusing lens, kinoform diffractive optic plate, and debris shield. A schematic of the layout is shown in Fig. 1. Other architectures were considered for NIF over the years, including Nova-like Master Oscillator–Power Amplifiers (MOPAs), two-pass systems, and others. CHAINOP and PROPSUITE have also been employed to look at the French Reverser/U-turn architecture.

For most NIF optimizations, this basic architecture was fixed, with the thickness of laser slabs and the number of slabs in each amplifier being the only layout parameters varied.

Laser Beam Size

CHAINOP and PROPSUITE use the same approach for determining the size of the laser beam that can be fit within a chain, starting with the hard (metal) aperture of the laser. This aperture can be square or rectangular, the current NIF design being 40 cm square. Determination of the size of beam that can fit within this aperture is dependent on the component spacings in the layout and necessary alignment tolerances. Component spacings determine the slight angle of the beam with respect to the chain axis necessary to pass through different pinholes on each pass in the spatial filters. As this angle increases, the size of beam that can fit within the hard aperture is reduced. This reduction in fill factor is known as vignetting. Separate pinholes for each pass are necessary primarily because ablated plasma on the edge of a pinhole can potentially block that pinhole to any subsequent passes of the beam. Pinhole separations are set by the more restrictive of either machining constraints in pinhole fabrication or sizes of optics that sit near the pinhole plane. These pinhole-plane optics include a mirror to inject the beam on the first pass in the transport spatial filter and a beam dump to catch that small fraction of the beam in the cavity not ejected by the Pockels switch on switch-out. The cavity filter pinhole spacings in NIF are set by machining constraints.

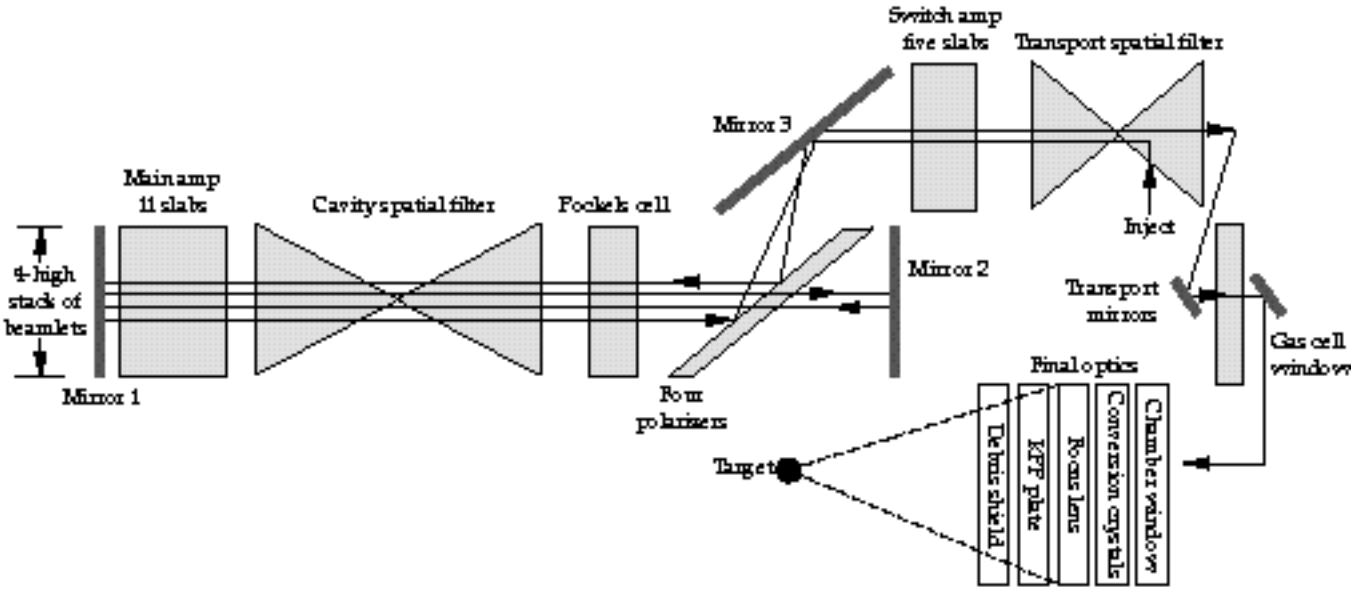

FIGURE 1. Schematic of NIF layout (not to scale). (70-00-1296-2745pb01)

An alignment tolerance of ~0.5 cm is included around the injected beam inside the hard aperture. This allows for beam or optics misalignment in the chain.

The beam intensity is nominally flat-topped, spatially, and tapers to zero at the edges. This taper is modeled with an error function, a shape somewhat more robust against diffraction ripples than some other shapes. Steeper edges increase the beam fill factor but simultaneously cause intensity spikes from edge diffraction and beam clipping at the pinholes. The taper distance of ~3 cm in the baseline design was determined as part of our optimization process.

The variables studied during NIF optimization affecting the beam size were hard aperture, number of laser slabs (which affects vignetting loss), and apodization tolerance.

Amplifier Gain and Pump-Induced Distortions

Gain of an amplifier slab depends on the stored energy density in the slab. This in turn depends on the pumping from flashlamps, Nd^{3+} concentration, slab thickness, and depumping due to Amplified Spontaneous Emission (ASE). The pumping term from flashlamps depends on the number of lamps per slab, their diameter, and electrical power pumped into the lamps. Both CHAINOP and PROPSUITE used a representation of the emission spectrum of flashlamps as a function of pumping power, absorption spectrum of laser glass, ASE depumping rate, and lamp-to-slab coupling efficiency. Coupling efficiency was based on 2-D ray-trace calculations, validated by measurements on Beamlet (a prototype of one NIF beamline in operation at LLNL), and is a function of reflector geometries, component spacings, slab thicknesses, and other parameters. Ray tracing was not included in the optimization directly. Rather, both codes used information from previous runs that covered the parameter space of interest.

Because PROPSUITE considers propagation effects on the beam spatial profile, we included pump-induced distortions in modeling. This distortion is the local steering imposed on the beam caused by the nonuniform heating of laser slabs and their subsequent deformation. Predicting this distortion requires three-dimensional (3-D) thermo-mechanical modeling of laser slabs under pumping conditions, giving time-dependent slab motion. These are lengthy computations, so prior to doing optimizations, a large number of cases were run covering ranges of the variables of interest. These results were stored in a database in which the code could interpolate to find distortions for any given combination of variable values.

Laser slab gain in both CHAINOP and PROPSUITE was modeled using the Frantz–Nodvick equations, which describe both the exponential growth of the beam fluence and the decay of the population inversion.[2] These equations are based on assumptions that (1) pumping to the upper laser level is negligible during the time the laser pulse extracts energy; (2) gain reduction due to spontaneous emission is negligible compared to stimulated emission during extraction; and (3) the cavity geometry is such that the laser pulse does not overlap on itself in space at any given time. With these approximations, extraction is only a function of fluence, and information regarding the detailed temporal shape of the pulse is not necessary. This allows us to reduce calculation time significantly by approximating a temporally shaped pulse with a temporally square pulse, while still accurately calculating the energetics of extraction and saturation. This is described in more detail below in the section entitled Temporal Pulse Shape.

The variables studied during NIF optimization affecting the gain and pump-induced distortion were thickness of the laser slabs, Nd^{3+} concentration in the laser glass, duration of the flashlamp pumping pulse, and number-per-slab and diameter of flashlamps.

Optics Losses and Aberrations

Optics losses in CHAINOP are simply subunity multipliers on beam energy and power. They include representations of laser glass surface scatter and bulk transmission, transmission of antireflection (AR) coated optics at 1ω and 3ω, bulk absorption of KDP/KD*P at 1ω and 3ω, and reflectivity of high-reflectivity (HR) coated mirrors. Bulk absorption of the fused silica is negligible for thicknesses of optics considered in NIF. Transmissivity and reflectivity of the polarizer were used to represent several effects, including coating reflectivity, BK7 substrate absorption, switching efficiency of the Pockels cell, and depolarization fraction of the beam due to stressed optics. Losses were not temporally dependent.

Losses were similarly included in PROPSUITE. In addition, by using a propagation code, it became possible to include spatially dependent representations of the surface finishes and bulk homogeneity properties of optics in the chain. This information was taken from a variety of measured parts chosen to reflect projected NIF manufacturing and finishing techniques. Because considerably fewer parts had been measured than were needed to represent an entire laser chain, it became necessary to synthesize the 100-plus necessary files. We did this by assuming that, for example, amplifier slabs will have the same finish as the measured parts, as quantified by the Power Spectral Density (PSD) curve.[3] We assumed, however, that phases of the Fourier components of the surface representations are randomized from one slab to another, resulting in less constructive addition of phase noise from slab to

slab. These aberration files are the dominant source for beam intensity and phase modulation in simulations.

During NIF optimization, losses and aberration files were not varied from a baseline set of assumptions.

Temporal Pulse Shape

NIF will use temporally shaped pulses for many experiments, including ICF ignition. To reduce calculational requirements, these pulses were represented in CHAINOP by an equivalent square pulse, having a power the same as the required peak power for a shaped pulse, and having the same total energy. This is a reasonable approximation because the Frantz–Nodvick equations depend only on beam energy (fluence), not the time history of power over the nanosecond durations of interest for NIF. Consequently, temporally square pulses of a given energy see the same overall gain as temporally shaped pulses of the same energy.

In CHAINOP, a square pulse is divided into two pieces: the main body and the tail end. Because gain in a laser slab is dependent on fluence history through the slab previous to that point in time, as the amplifier saturates, the tail end of the pulse sees a significantly lower gain than the front. A pulse that will be temporally flat at the end of the chain must, therefore, be temporally increasing at injection. Beam filamentation —catastrophic self-focusing due to the effects of a nonlinear refractive index—depends on the instantaneous beam irradiance. Consequently, monitoring the power of the pulse's tail in CHAINOP allows us to track and limit the tendency of the beam to filament.

In contrast, PROPSUITE is capable of tracking the beam shape at an arbitrary number of times. We used eight time steps to represent the temporally shaped indirect-drive ICF ignition pulse (See Fig. 2). For this pulse, the maximum effect from the nonlinear index of refraction of the optics occurs somewhat before the end of the pulse and at a time that varies with position in the chain, justifying this more detailed modeling.

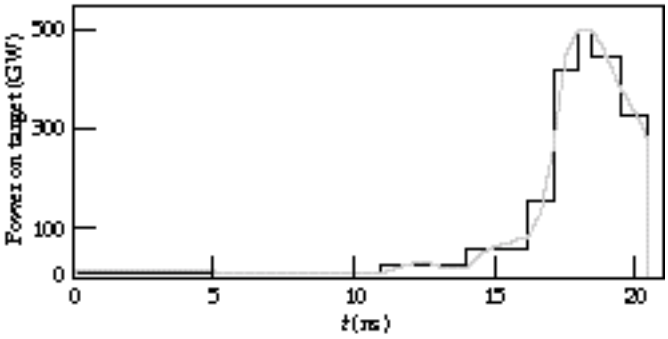

FIGURE 2. Temporally shaped pulse used for optimizations. The gray line is the actual pulse shape needed on target. The black line is the eight-step approximation to this shape used in modeling. (70-00-1296-2748pb01)

Frequency Conversion

The frequency conversion efficiency (3ω irradiance / 1ω irradiance) can be a function of space and time given temporal or spatial variations of the beam. Since neither temporal nor spatial modulation is calculated in CHAINOP, both peak-power conversion efficiency (averaged across the beam aperture) and energy conversion efficiency for temporally shaped pulses needed for NIF are code inputs. Power conversion efficiency is used to determine the filamentation threat to the 3ω optics and peak power on target. Energy conversion efficiency is used to determine energy on target and damage threat calculations.

In PROPSUITE both temporal and spatial intensity information are available. Consequently, we use three codes to calculate conversion efficiency directly. A plane-wave / temporally varying field code is used to predict the 1ω temporal shape that would give the required 3ω pulse on target for any given converter design. The spatially and temporally resolved 1ω pulse at the converter with this temporal shape is then generated with PROP92. The conversion of this beam is then calculated with a full spatio-temporal code that includes the capability of having converter-crystal surface-finish noise added to the beam phase. See the article "Frequency-Conversion Modeling" on p. 199 of this *Quarterly* for a discussion of these frequency conversion codes.

For NIF optimizations, thicknesses of the type I doubler (KDP) and type II tripler (KD*P) and the detuning angle of the doubler crystal were varied independently. To allow for difficulties in aligning crystals with respect to the beam, and for other problems, conversion calculations for the 1ω beam were actually done at three detuning angles of the doubler, a variable center angle and ±30 μrad around that. A weighted average of the three was then used to determine energy on target.

Beam Modulation

Since CHAINOP is zero-dimensional, the spatial intensity / fluence modulation of the beam cannot be directly calculated. This information, however, is necessary for evaluating damage to optics from high fluences. Accordingly, a semi-empirical rule was implemented that the peak-to-mean modulation in fluence (φ) is given by

$$\varphi_{peak} = \varphi_{avg}(1.3 + 0.1e^{\Delta B}) \quad . \tag{1}$$

ΔB is the B integral accumulated by the beam between pinhole passes (using approximately the mean spatial intensity). The B integral, in general, is defined as

$$B(z) = \frac{2\pi}{\lambda}\gamma_{nl}\int_0^z I(z')dz' = \frac{80\pi^2 n_2}{\lambda cn}\int_0^z I(z')dz' \quad , \tag{2}$$

and is the nonlinear phase retardation (in radians) experienced by a beam of intensity I (W/m²) and wavelength λ (m) in traversing a medium of thickness z (m) with refractive index n_0 and nonlinear index n_2 (from $n = n_0 + n_2 \hat{E}\hat{E}^2$). (The nonlinear index may also be a function of z in nonisotropic materials.) Bespalov–Talanov theory predicts[4] that small-scale spatial ripples in the beam will grow nonlinearly, with the most unstable perturbations growing as $e^{\Delta B}$. This, combined with the fact that beam modulation past a scatterer in the absence of nonlinear effects is no more than 1.4:1, and that in large laser systems it has been experimentally observed that beam modulation is ~2:1 after a ΔB of 2 radians, leads to Eq. (1).

In PROPSUITE, beam modulation is calculated directly, including diffraction and nonlinear effects.

Optics Damage Thresholds from High Fluence

High laser fluences on optics over nanosecond-pulse durations of interest for NIF can cause surface or bulk damage sites. This damage is seen as pits in the surface or coating of an optic, or damage spots in bulk. The fluence threshold when this damage occurs is experimentally found to be a function of the pulse duration. For temporally Gaussian pulses, the experimental database is well matched by the functional form fluence damage threshold (J/cm²) $= a\tau^b$, where a and b are experimentally determined and material-dependent constants and τ (ns) is the full-width-half-maximum of the pulse.

NIF pulses will typically range in temporal shape, and may be significantly non-Gaussian. To account for the effect of arbitrarily shaped pulses on damage threshold, we used a damage diffusion model[5] to determine the damage thresholds. Because the ICF indirect-drive ignition pulse shape used for optimizations is more square than Gaussian, the result is a 5 to 20% reduction in damage thresholds from those for a Gaussian pulse for various materials and coatings. This diffusion model adjustment of damage threshold has not yet been experimentally tested.

The 3ω optics (tripler, final focus lens, kinoform phase plate, and debris shield) are exposed to 1, 2, and 3ω light simultaneously. In general, damage fluence limits at the longer wavelengths are significantly higher than at shorter wavelengths, but we have found no additional information regarding multiwavelength interactions. Accordingly, we have based our damage limits on the 3ω fluence only, with a damage threshold decreased by 10% from its value under monochromatic illumination. This assumption remains an area of uncertainty in our modeling, since we predict system performance to be limited by 3ω damage in the tripler.

Both CHAINOP and PROPSUITE used the same damage thresholds.

Filamentation Risk

Filamentation damage (or "angel hair" damage) is seen when the nonlinear index of refraction of an optical material causes self-focusing of small sections of a high-intensity beam. These sections may start out as small-intensity (a few percent above average) ripples or bumps of a few millimeters in size on the beam. As they self-focus, they decrease in size to hundreds of micrometers in diameter or smaller, and increase to very high irradiances (>100 GW/cm²). At some point the irradiance is high enough to damage the material and leave a visible track.

This filamentation growth is a complex function of the beam irradiance, nonlinear index of the material, and size and shape of the irradiance perturbation that self-focuses. It is also an inherently three-dimensional problem, making modeling calculationally intensive, although possible with present tools.[6] However, it is in fact not necessary to model growth of a filamenting bump to the point of damage in order to design NIF because we desire to be far from this limit. Consequently, a more heuristic approach was used in the optimizations, based on the experimental observation that *if small-scale beam noise is present*, large lasers tend to begin to break up into filaments after the beam has experienced a B integral of ~2 radians. This deserves a little more explanation.

The presence of small-scale beam noise is important because bumps of small size (a few millimeters) grow significantly faster than large-scale modulations. This difference in growth rate has been understood for many years.[4]

Elimination of these small-scale noise ripples is one of the purposes of spatial filters in large laser systems. Each laser chain in NIF will have two such spatial filters, one in the cavity, and one for transport to the target, as shown in Fig. 1. At the focal plane of each filter there is a pinhole through which the beam passes. The pinhole size, however, is such that only spatial frequencies present in the beam (in either phase or intensity) of less than a certain value can pass through the pinhole. Power at frequencies higher than this cut-off value is dumped on the surface around the hole. For NIF this cut-off half-angle is 100 μrad, corresponding to spatial ripple wavelengths of 1 cm. The net effect of this spatial filtering is that perturbations with transverse wavelengths less than 1 cm are clipped off at the pinholes. Consequently they cannot grow for more than the B integral they experience between pinhole passes.

In the chain design, then, we can limit B integral seen by the laser between pinholes to a certain value as a surrogate for calculating the beginnings of filament growth. For NIF optimizations, the between-pinhole B integral limit (also known as ΔB) was set to 2.2 rad. This constraint was used in both CHAINOP and PROPSUITE.

Optimization Variables and Constraints

Techniques used in laser system optimizations are discussed in the article "Laser Optimization Techniques" on p. 192 of this *Quarterly*. Our codes were implemented with simplex, parabolic interpolation, and gradient-search routines. The figure of merit that the optimizer attempted to minimize by changing variables, and, at the same time, honoring several constraints, was system cost divided by energy on target. These variables included hard aperture, injection fluence, number of laser slabs in each amplifier, laser slab thickness, Nd^{3+} concentration in laser slabs, flashlamp packing fraction, and apodization edge width. In using a gradient-search optimization technique, we took care in code development that there were no discontinuities in value or 1st derivative of figure of merit or constraints, as a function of these variables. Constraints that the optimizer used included energy on target, fluence damage to components, and ΔB limits.

Cost Algorithms

We have used different levels of cost modeling during development of the NIF design. Those used in later calculations included cost scaling rules, as a function of the optimization variables, for ~30 nonoptical subsystems of NIF, and ~140 optical cost categories. These rules were built on a cost database in the NIF *Conceptual Design Report*, and were provided by the NIF engineering team.

NIF Optimizations

Optimization studies with CHAINOP resulted in the CDR design. Similar studies with PROPSUITE resulted in the Advanced Conceptual Design (ACD) layout. The current NIF design (Title I, or preliminary design) has changed little overall since the ACD, so, for purposes of this discussion, the two will be considered equivalent. In this section, we describe the ACD optimization process and compare those results to the conceptual design.

During the ACD, we used three optimization approaches. In two approaches, optimization routines were used to find the least expensive design that meets functional requirements. These methods yielded a high-quality answer, but gave limited insight into tradeoffs in design options. The third method was to evaluate a large number of separate configurations manually, covering some range of parameter space on a uniform grid, and to graphically sort data according to design criteria. This approach enables greater understanding of the design space, but does not offer assurance that the best possible design has been included on the grid. The combination of this and the optimization approaches gives confidence that available design space has been explored thoroughly. Results from these methods are discussed below.

Studies Using Optimizing Software

By decision of the NIF project, some parameters were not varied during the final ACD optimization. Specifically, after some preliminary results, the optical clear aperture of the laser was fixed at 40 cm, laser slab thickness at 4.1 cm, and transport spatial filter length at ≥60 m. The optical clear aperture and slab thickness values were both larger than those found by the unbiased application of our optimization process. They were chosen to increase the 1ω performance margin (larger stored energy) and to decrease the risk of beam filamentation (lower between-pinhole B integral). The minimum transport spatial filter length was set primarily from concerns about pinhole blowoff in the last beam pass through the filter. (A shorter filter would proportionately decrease the pinhole size necessary to filter features of given scale. This would both increase the irradiance on the pinhole edge and decrease the closure distance.) The detailed physics of pinhole closure is currently poorly understood, making a more conservative choice advisable. Other values for these parameters would have saved additional money, but at the cost of decreased experimental flexibility and increased technical risk.

Given these constraints, we performed an optimization to determine the least expensive design overall, and also the least expensive design with no amplifier adjacent to the Pockels cell switch (the current design). These results are shown in Table 1. They are compared to the design arrived at during the CDR using CHAINOP (with small changes to allow better comparison between cases). The expected reduction in NIF cost by moving from the CDR design to the other two designs, as determined from the cost algorithms, is also included.

It can be seen in this table that the least expensive design (with the above constraints) is a 9/5/3 layout (the layouts are referred to by the number of slabs in each of the three amplifiers); the 9/5/3 layout is approximately $10M less expensive than the CDR design, but has a somewhat higher peak ΔB. The optimal design without a switch amp is the 11/0/7 design, with a cost approximately $4M less than that in the CDR. The $4M difference between the CDR and Title I designs is a small fraction of the total cost (~0.5%). The fact that CHAINOP, used for the CDR optimization, would suggest a design so close in cost (fractionally) to the Title I design, developed by PROPSUITE, gives confidence that CHAINOP provides accurate answers in spite of its relative simplicity.

Although we predict that both the 9/5/3 and 11/0/7 designs will meet the performance criteria that were set, there are a number of operational reasons to

TABLE 1. Comparison of the CDR design (9/5/5) and ACD optimized cases (9/5/3 and 11/0/7).

Parameter	Modified* CDR with new converter	Least expensive, constrained,** optimized	11/0/7, constrained,** optimized
Hard aperture (cm)	40	40	40
Main amp # slabs	9	9	11
Switch amp # slabs	5	5	0
Boost amp # slabs	5	3	7
Slab thickness (cm)	3.36	4.1	4.1
Injection energy (J) ***	3.2	1.2	2.2
Doubler thickness (cm)	1.35	1.36	1.36
Tripler thickness (cm)	1	1	1
Doubler detuning angle (μrad)	200	200	200
Flashlamp pump pulse length (μsec)	360	390	360
Nd^{3+} concentration ($10^{20}/cm^3$)	4.18	3.87	3.59
Transport spatial filter length (m)	65	60	60
Last two ΔB's (radian) (at beam center)	1.3/2.05	1.7/2.2	1.7/1.8
Conversion efficiency (energy/peak power) (%)****	60/80	60/79	60/79
Capacitor bank (MJ)	373	347	353
Total length (M1 to L4) (m)	135	128	123
Energy (MJ)/peak power (TW) in 600 μm LEH ***	2.2/612	2.2/608	2.2/607
Cost relative to CDR (M$)	0	–10	–4

* Changes from CDR: 1 × 1-cm pinhole spacing in cavity, new gain/loss assumptions, 4 × 2 amplifiers, 3.5-cm spatial filter lenses, and addition of target chamber window. These assumptions were common to the other two designs.

** Constraints: 192 beamlets, 40-cm hard aperture, 60-m transport spatial filter, 4.1-cm thick slabs.

*** Before spatial preshaping for gain roll-off across the aperture.

**** Averaged over ±30-μrad doubler detuning angle.

prefer the 11/0/7 layout. Most apparently, from Table 1, it has a noticeably lower maximum ΔB, yielding some additional safety against nonlinear effects. Equally compelling, the similarity to the 11/0/5 design for Beamlet (the operational prototype NIF beamline) makes operational experience gained on that machine more relevant and will likely decrease prototyping costs and/or risks. The absence of a switch amplifier also both obviates concerns about the effects of magnetic fields on the operation of the plasma electrode Pockels cell and greatly simplifies the alignment procedure. There are also some design cost savings in not having a switch amplifier included in the analysis. For all of these reasons, the NIF project chose the 11/0/7 design as its baseline for the ACD.

The shift to the current NIF 11/0/5 design was done as part of the Title I preliminary design activities. It entailed slightly different assumptions than used here, with somewhat increased risk, in order to realize a lower cost.

Studies Using Parameter Scan Optimization

In addition to optimizations, we used PROPSUITE to perform parameter scans of the cost and performance of a large number of designs (~50,000) in the vicinity of the optimized NIF design. Parameters varied for this study were slab thickness, number of slabs in each amplifier, Nd^{3+} concentration, and flashlamp pump pulse length. This was done for 26 different combinations of the slab counts in the three amplifiers ($X/Y/Z$).

Figure 3 shows the results for cases with no switch amp ($Y = 0$). Cost is plotted vs output energy, with the several points for each slab combination being different combinations of Nd^{3+} concentration, slab thickness, and pump pulse length. The temporally shaped ICF indirect-drive pulse from Fig. 2 was used in each case. All other parameters were the same between designs (including a 40-cm hard aperture). It can be seen that there is some grouping of the various designs, with the least expensive designs being 11/0/7, 11/0/5, and 13/0/5. Some designs with tighter grouping show a greater robustness to parameter changes than others (e.g., 11/0/7 vs 9/0/7). This is desirable in a system. The designs at higher cost are typically characterized by thin slabs and large injection energies. For instance, the large number of thin slabs incurs a higher slab finishing cost and larger capacitor bank. Too few slabs in the boost amplifier, which is only double-passed, requires a larger injection energy and a more expensive front end.

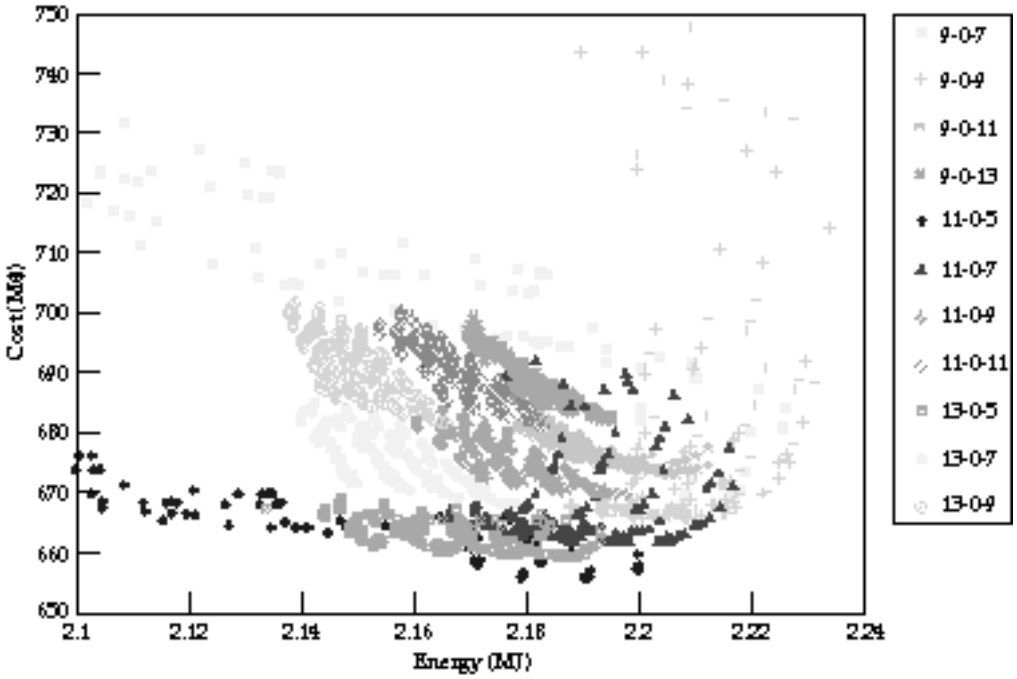

FIGURE 3. Full parameter scan results for cases with no switch amp (*X*-0-*Z*).
(70-00-1296-2779pb01)

All the designs in Fig. 3 produce slightly different energies and peak powers on target. In order to simplify comparisons, it is possible to normalize all the designs to have 2.2 MJ/600 TW on target by projecting small changes in hard aperture and by scaling cost accordingly. This scaling of the aperture leaves the B integral unchanged for each design by increasing/decreasing the beam area. For example, if a design produced 2.15 MJ and 580 TW, it would be scaled to a 40.7-cm hard aperture, giving 600 TW and 2.22 MJ. The cost was scaled as (hard aperture)$^{0.8046}$, based on studies done previously with CHAINOP.

With this normalization, we plotted the various cases against maximum ΔB in the system (Fig. 4). In this display it can be seen that the 11/0/7, 11/0/5, and 13/0/5 designs are still least expensive, but that they have substantial differences in maximum ΔB, with the 11/0/7 design displaying the lowest ΔB.

In Figs. 5 through 7, successive filters are applied to the data in Fig. 4. A large injection energy requirement typically indicates a design where the modeled pulse is close to the extraction limit of the system. This leaves system performance very vulnerable to small uncertainties in losses or gains. Designs with injection energies of a few joules or less are safer; Fig. 5 shows cases that have injection energies ≤5 J. The most attractive designs are still the 11/0/7, 11/0/5, and 13/0/5, although many of the 11/0/5 designs with thinner slabs were eliminated.

Designs with overly large capacitor banks (and correspondingly small injection energies) are simply uneconomical, as more energy is stored than necessary. These can also be thought of as designs with too many slabs. Figure 6 shows cases filtered for capacitor bank size ≤400 MJ. The 11/0/7, 11/0/5, and 13/0/5 designs are still most attractive. Note that the optimum number of slabs is 16 to 18.

Results of a final filtering requiring laser slab thicknesses ≤4.2 cm are shown in Fig. 7, reflecting fabrication uncertainties with thick slabs. This filter removes all the remaining 11/0/5 designs, leaving the 11/0/7 and the 13/0/5 designs as having comparable minimum costs. The 11/0/7 designs, however, have lower B integral. Clearly, given these criteria, the best design with no switch amplifier is the 11/0/7.

We can compare these results to cases with a switch amplifier. In Fig. 8, a subset of cases is shown having slab counts of *X*/0/*Z*, *X*/3/*Z*, and *X*/5/*Z*, with all the aforementioned filters applied. It is first useful to note that there are a large number of cases with approximately the same cost ($660–$670M), but some difference in B integral. Given that, however, it can be seen that the 9/5/3 design is the least expensive design, with the 9/3/5 and 11/0/7 close behind. These results agree with those cases found by the optimizer, and shown in Table 1: the 9/5/3 design is the least expensive, but has a somewhat higher B integral than the 11/0/7 (2.1 to 2.2 vs 1.7 to 1.8 rad). The design chosen by the project is circled on Figs. 7 and 8. The 11/0/7 designs that have slightly lower cost and B integral than the one selected have thinner slabs than the prescribed 4.1 cm.

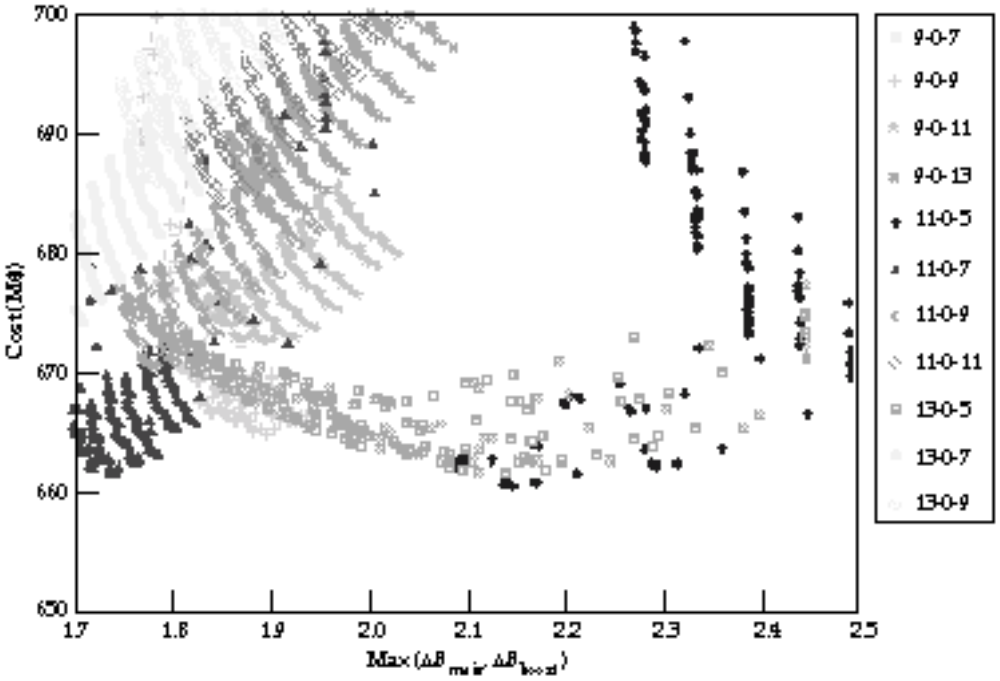

FIGURE 4. Data from Fig. 3 (*X*-0-*Z* parameter scan), aperture-scaled to 2.2MJ, and compared on the basis of maximum B integral.
(70-00-1296-2780pb01)

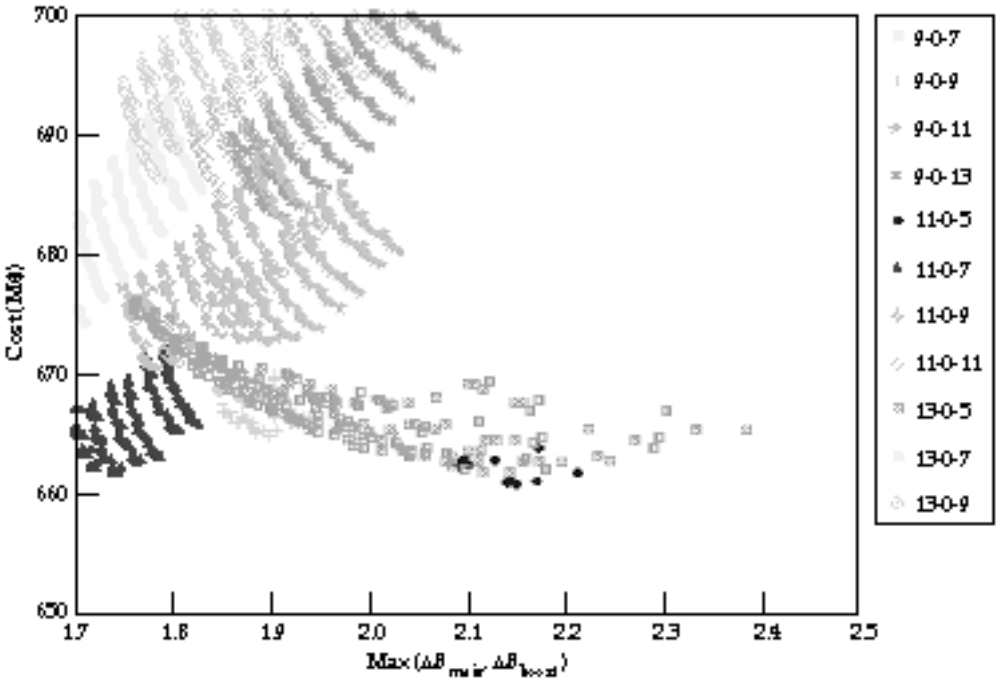

FIGURE 5. Data from Fig. 4 (*X*-0-*Z* parameter scan), clipped at injection energy ≤5 J.
(70-00-1296-2781pb01)

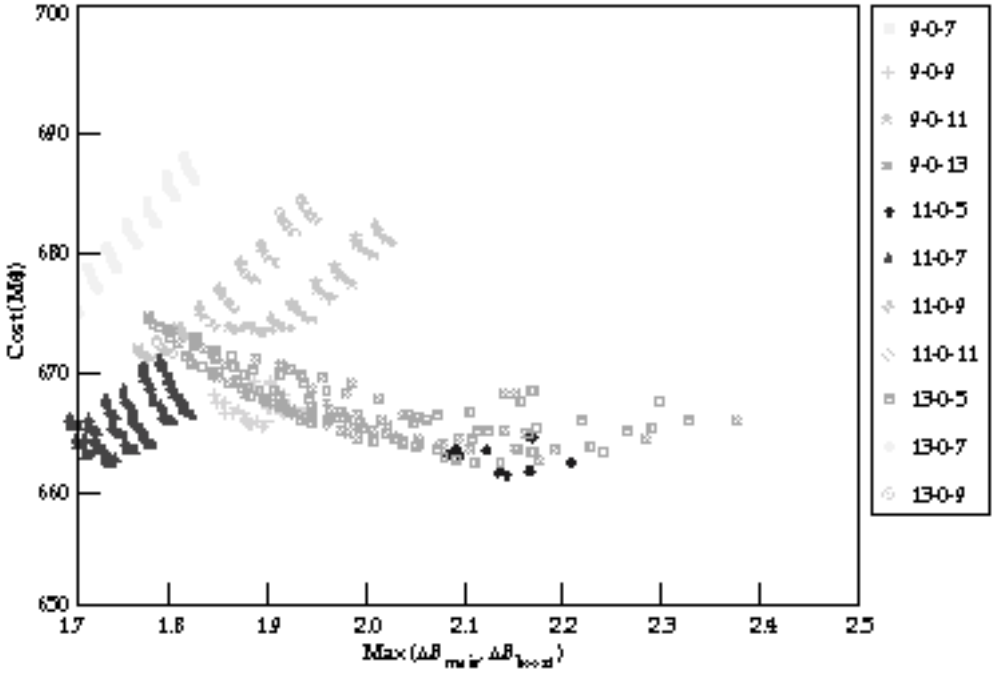

FIGURE 6. Data from Fig. 5 (*X*-0-*Z* parameter scan), clipped at injection energy ≤5 J and capacitor bank size ≤400 MJ. (70-00-1296-2782pb01)

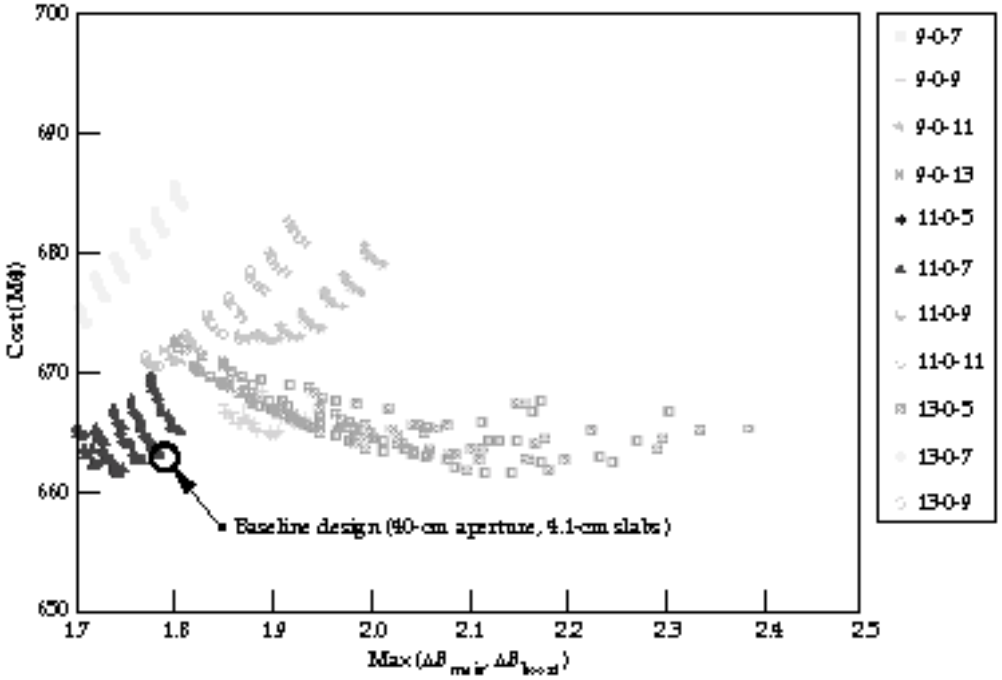

FIGURE 7. Data from Fig. 6 (*X*-0-*Z* parameter scan), clipped at injection energy ≤5 J, capacitor bank size ≤400 MJ, and laser slab thickness ≤4.2 cm. (70-00-1296-2783pb01)

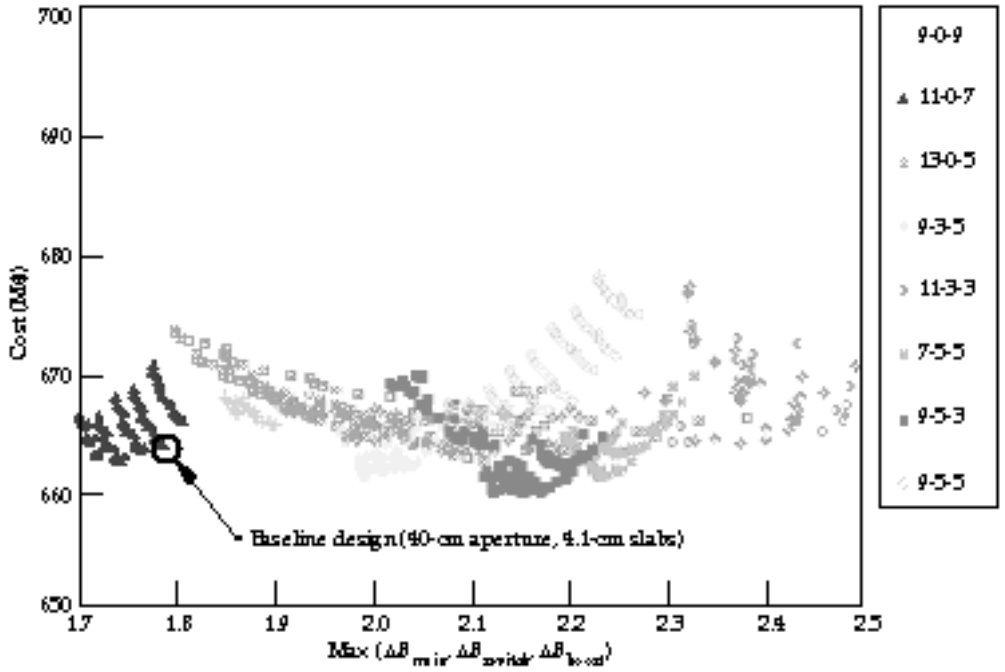

FIGURE 8. Comparison of subset of parameter scans for *X*-0-*Z* and *X*-*Y*-*Z* with injection energy ≤5 J, capacitor bank size ≤400 MJ, and laser slab thickness ≤4.2cm. (70-00-1296-2784pb01)

Summary

The NIF design is based on cost/performance optimization studies. These studies were done in two phases, and with two codes: a zero-dimensional performance model (CHAINOP) for the *Conceptual Design Report*, and a one-dimensional propagation model (PROPSUITE) for the Advanced Conceptual Design (ACD). These ACD results using PROPSUITE led to 11/0/7 design of Table I. The current Title I 11/0/5 design activity was a result of further cost and risk studies. These efforts give us confidence that we are designing the (approximately) least expensive system meeting the functional requirements.

Acknowledgments

The authors would like to thank Mark Rotter (LLNL) for assistance in setting up and running the NIKE3D and TOPAZ codes used in thermo-mechanical modeling of laser slabs.

Notes and References

1. *National Ignition Facility Conceptual Design Report,* UCRL-PROP-117093, Lawrence Livermore National Laboratory, Livermore, CA, 94550 (1994).
2. L. M. Frantz and J. S. Nodvik, *J. Appl. Phys.*, **34**(8), 2346 (1963).
3. J. Lawson, *Proc. Conf. on Optical Manufacturing and Testing,* SPIE Vol. 2536, pg.38 (July 1995).
4. J. Trenholme, *LLNL Laser Program Annual Report*, UCRL-50021-74, pg. 179 (1974).
5. J. Trenholme, personal communication.
6. W. Williams, K. Manes, J. Hunt, P. Renard, et al., *ICF Quarterly Report*, **6**(1), 7, Lawrence Livermore National Laboratory, Livermore, CA UCRL-LR-105821-96-1 (1995).